Visualization and Collaboration over Optical Networks

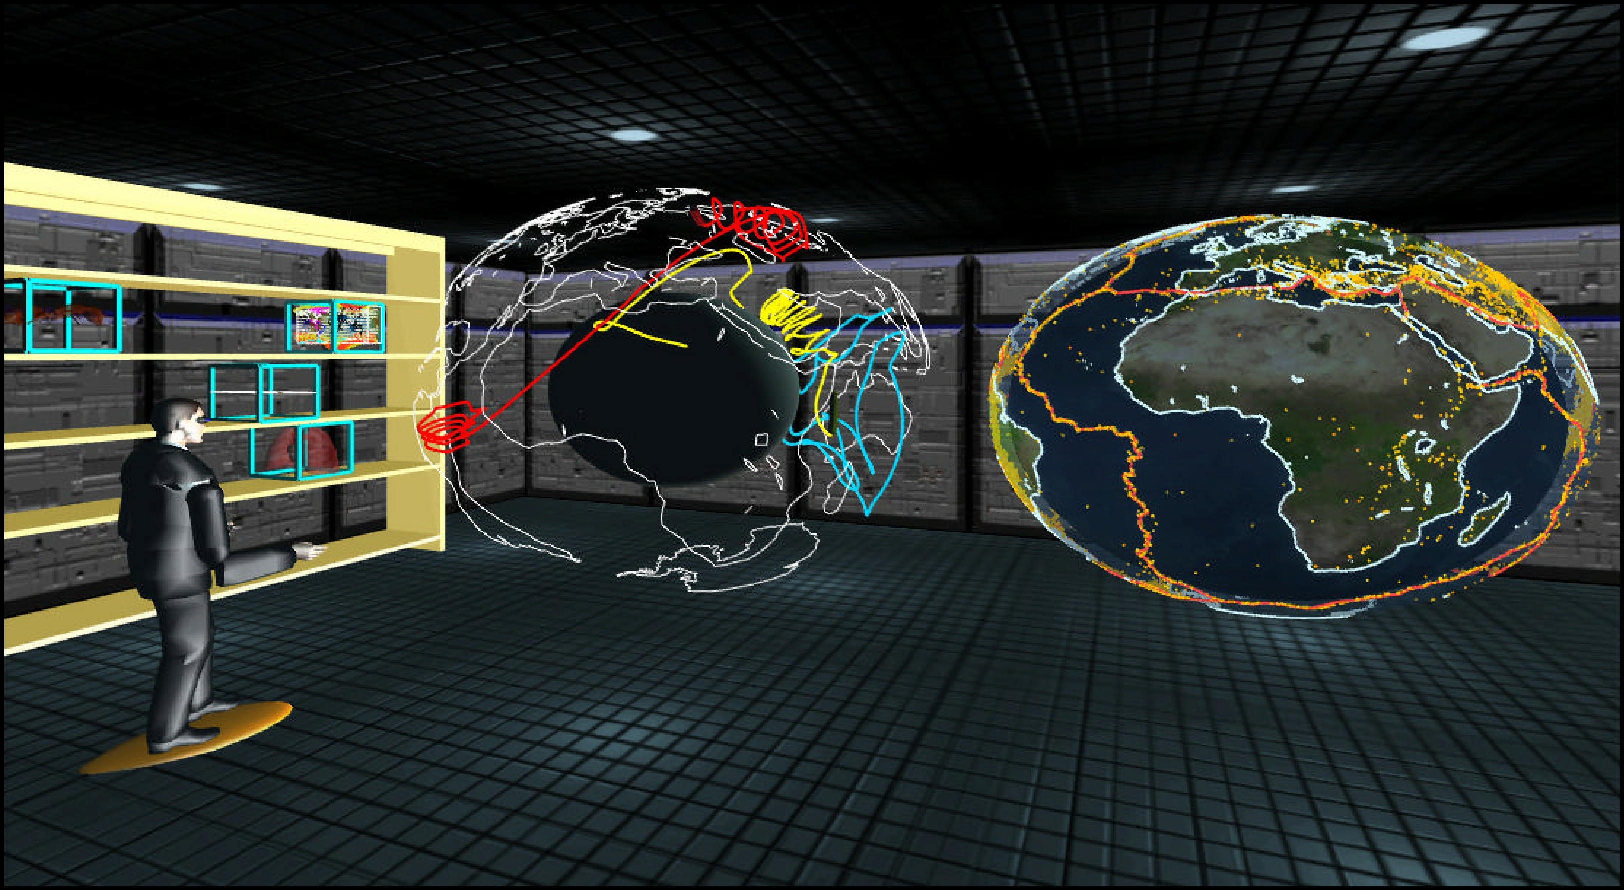

Jason Leigh, Robert Grossman, Oliver Yu, Alan Verlo, Linda Winkler,
Maxine Brown, Dan Sandin, Thomas A. DeFanti

Electronic Visualization Laboratory (EVL) University of Illinois at Chicago

- State of Illinois has paid for 20 year lease of fiber between UIC, U Chicago, NWU, ANL, NCSA
- 10Gbps between campuses
- STARLIGHT

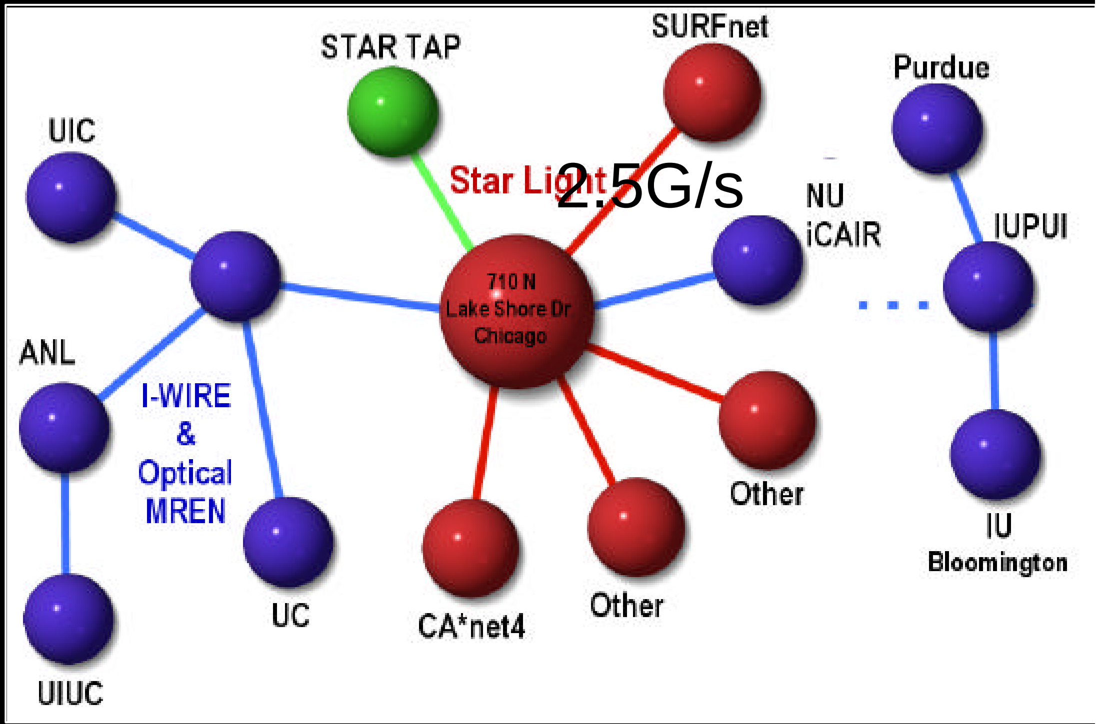

TeraNode

- Basic TeraNode
  - "Hard to walk and chew gum at the same time"
  - 20 processors, each with:
    - GHz processor
    - 2-4G RAM/processor
    - 20G disk/processor
    - Gigabit ethernet card
  - Optical networks become the system bus for these nodes
- TeraNode-D (TND): Data
  - Multi-terabyte RAID disks
- TeraNode-C (TNC): Computation
  - Large number of processors and memory
- TeraNode-V (TNC): Visualization
  - High end graphics cards and lots of memory
  - Various configuration of displays from passive stereo walls, to CAVEs to tiled displays

Starting Simple: Identifying bottlenecks
Stereoscopic Graphics Streaming

Video server L (TND) —500Mbps→ Video client L (TNV)

Video server R → Video client R

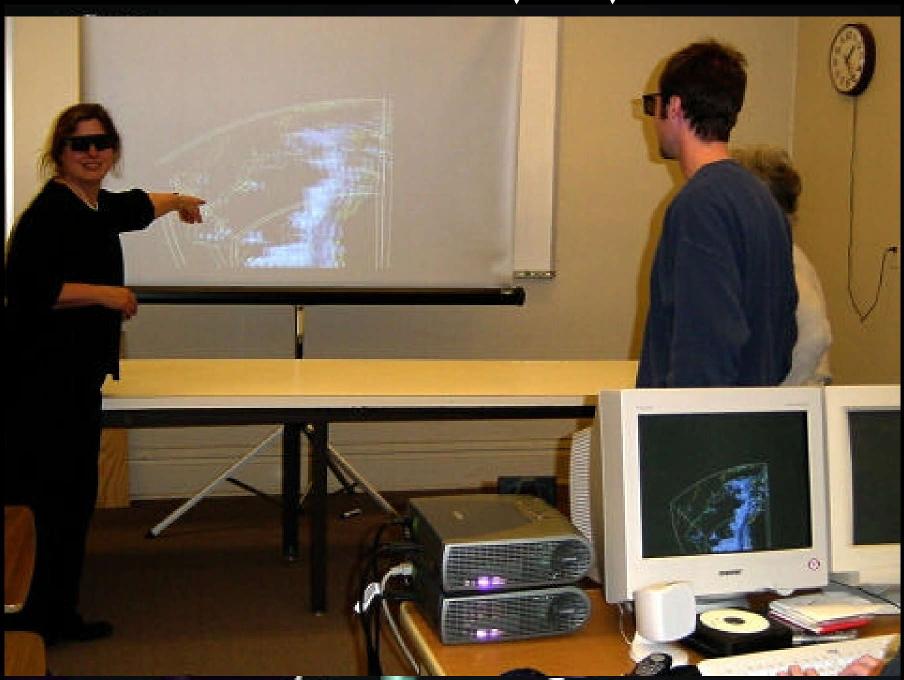

1024x768 30fps
uncompressed 24 bit graphics

- Bottlenecks- GigE cards, OS versions, Graphics cards & drivers
- For collaboration how do you support multicast over optical networks?
- Stream Zbuffer?

TeraVision: Tiled Cluster to Cluster Graphics Streaming

TND → TNV tiled display

- EVL is building a tiled LCD autostereo (Varrier) display wall
- Uncompressed 1024x768, 24bit, 30fps, 5x4 tile ~ 10Gbps
- Will require GMPLS provisioning of light paths to guarantee QoS
- Could store movie on local disk and then play it back but:
  - 78 GBytes/s of animation
  - 1 Terabyte can hold 14 minutes of animation

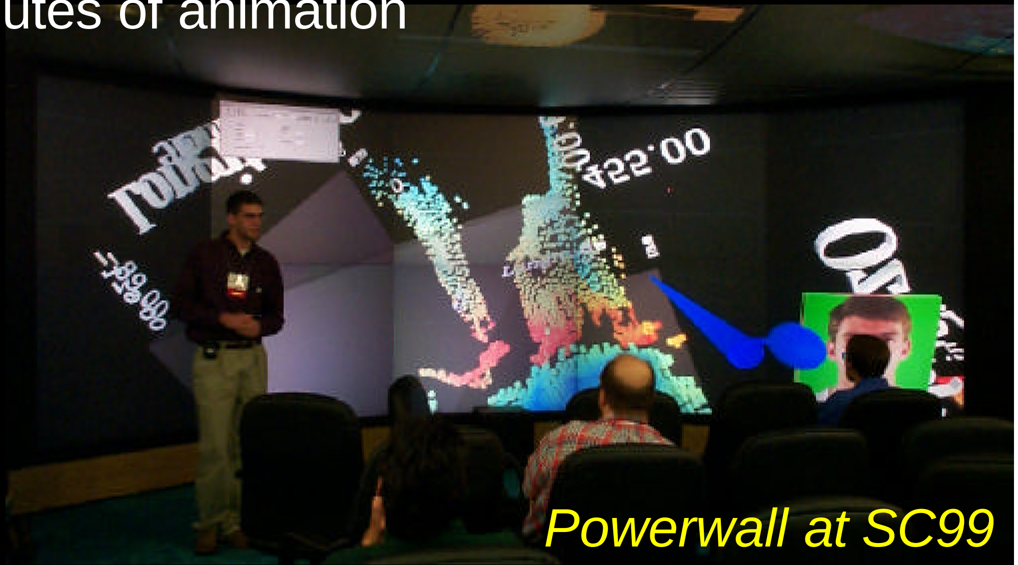

Powerwall at SC99

Streaming Visualization Systems Worth Testing Over Gigabit Networks

- WireGL/Chromium – streaming OpenGL draw commands
- VNC – streaming desktop graphics
- Pipelined VTK – streaming MPI messages and data for parallel rendering
- Worthwhile to monitor them and see what kind of performance improvement one gets, if any, and tune or parallelize underlying algorithms if needed.

Real Time Tele-Immersive Visualization of Massive Data Sets

- Tele-Immersion is collaborative virtual reality augmented with large scale computation and data access.
- Stereo and head tracking (VR) shown to increase amount of info that a viewer can interpret [Ware96]
- Collaboration useful for partitioning tasks that are too large for single person to handle
- Large Scale Data Viz in Tele-Immersion is difficult because of Real Time requirements of VR
- SGI Onyxs are not giving us quantum leaps in graphics power
- PC graphics card look promising but benchmarks show they they are fill limited[Pape01]

Real Time Tele-Immersive Visualization of Massive Data Sets

- Problem in Tele-Immersion is that Visualizations may be:
  - Small enough to display in real time but there are too many frames to fit in memory
  - Visualizations are too large to display in real time
  - Both!
  - And want to be able to compare visualizations side by side

Cached visualizations

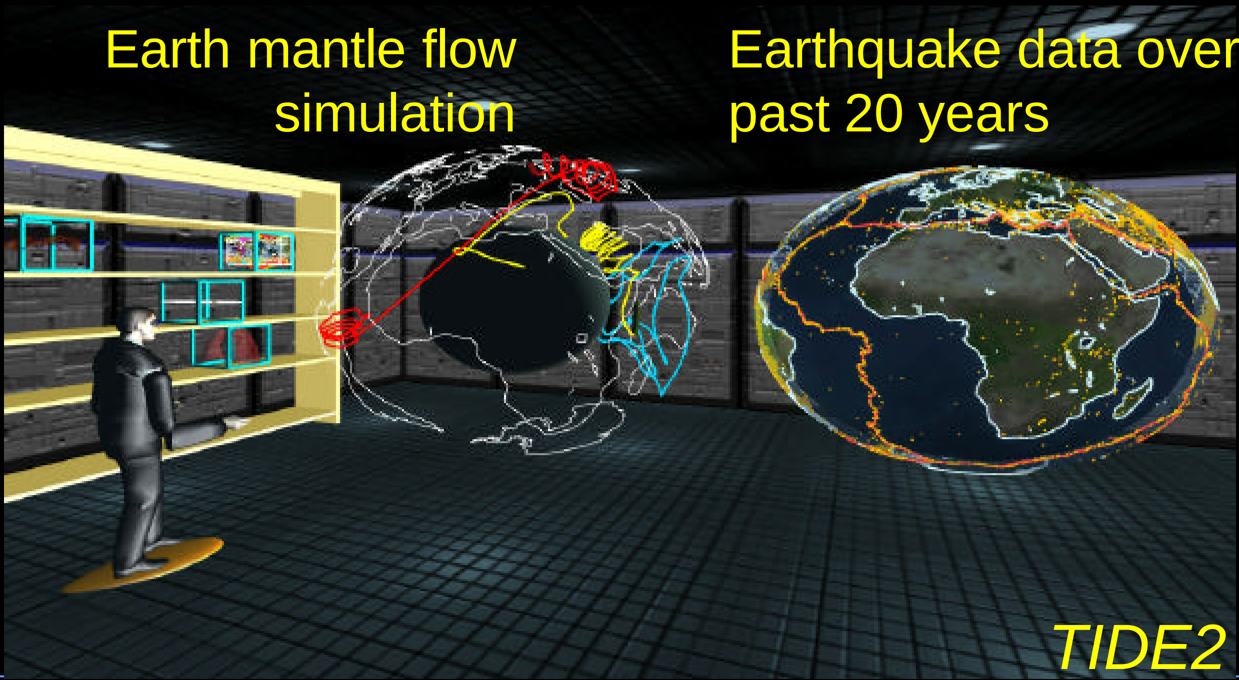

Real Time Tele-Immersive Visualization of Massive Data Sets

- Stream visualizations from TNCs as pre-rendered view-dependent rasters to immersive environments as billboards with alpha channel to eliminate background – Virtual Sprite
- Use WireGL/Chromium
- Develop memory management algorithms to page visualizations in and out from neighboring GigE-connected TNDs- what we need is Network RAM

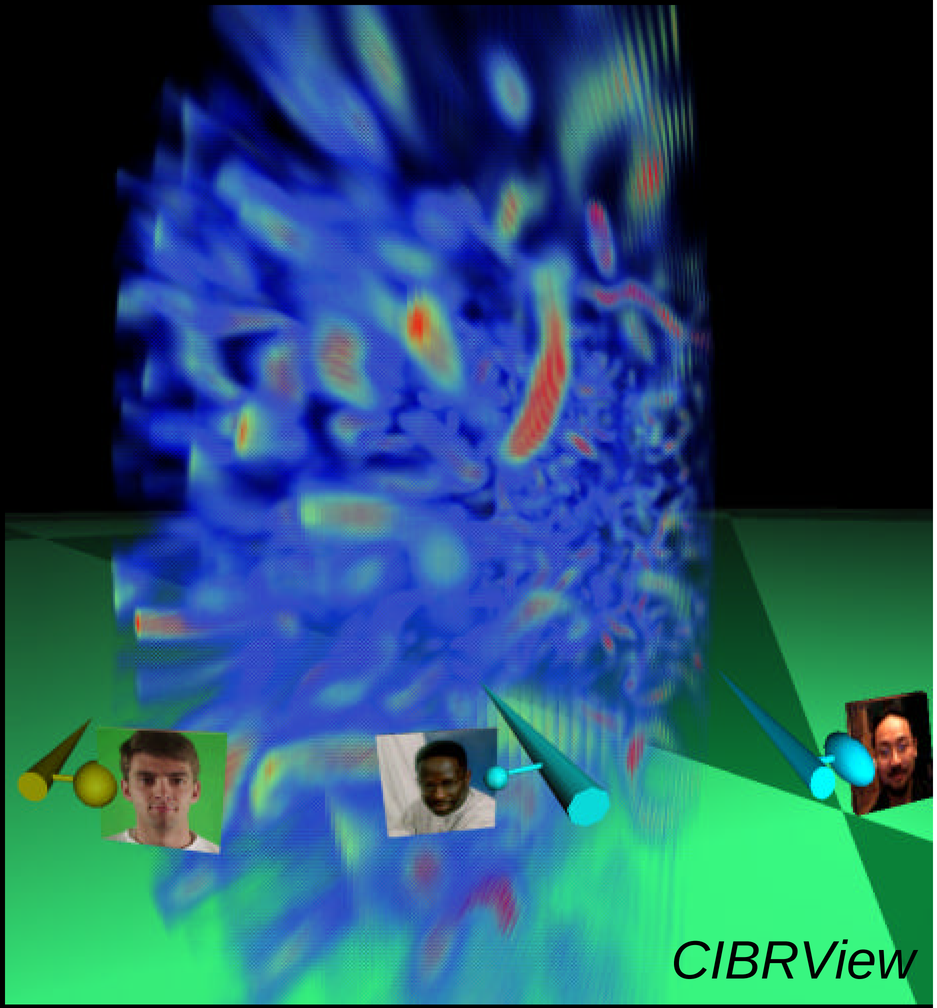

Interactive Visual Tera Mining
(Visual Data Mining of Tera Byte Data Sets)

Chicago

Data Mining Servers (TND-DSTP)

Amsterdam

Data Mining Servers (TND-DSTP)

NWU, NCSA, ANL etc…

Data Mining Servers (TND-DSTP)

Parallel Data Mining Correlation (TNC)

Parallel Visualization (TNV)

Tera Map

Tera Snap

Tera Snap

Tera Snap

– Problem is to touch a Terabyte of data interactively and to visualize it
– 100M/s – 24 hours to access 1 Terabyte of data
– 500M/s – 4.8 hours using a single PC
– 10G/s – 14.4 minutes using 20 node PC cluster
– Need to parallelize data access and rendering

Adaptive Networking : The Motivation

- Bandwidth is becoming increasingly available. E.g. Japan is aggressively targeting 100Mbps fiber to the homes (expect all homes by 2005).
- Networking QoS is still under research and difficult to deploy and use. It is not useable at a “flip of a switch.”
- Regarding QoS, Networking engineers think only bandwidth is important- not latency and jitter.
- Applications today are not ready to use the extra bandwidth even if available.
- Application developers have to be increasingly network savvy in order to be able to map application requirements to networking services.
- Make networking easier for the average application developer.

QUANTA : Quality of Service and Adaptive Networking Toolkit

1. Network QoS scheduling, provisioning, and at a lower level, control (to maintain QoS levels) – Generalized Multi Protocol Label/Lambda Switching
2. Socket tuning or intelligent transport protocol selection- on optical networks the edges need to be smarter since the core is just switching photons.
   1. Parallel TCP
   2. Forward Error Correction
   3. Reliable Blast UDP, etc..
3. Monitoring so that you don't ask for more bandwidth than you can actually use
4. An adaptive networking system/framework to take a user/application-specified QoS requirement and use 1,2,3 to satisfy the requirement

OptiMem : Network Memory over Optical Networks

- Latency to page to disk (~13ms- Fast Wide SCSI2) is greater than latency to access memory remotely (from EVL to ANL: 1.5 ms)
- Disk bandwidth ~ 250Mbps
- Create massive memory by networking TeraNodes over optical networks
- => 4 Levels of Memory: Local, Cluster, NetworkRAM, Disk
- For higher latency networks use pre-fetching to compensate since additional bandwidth is available.
- Provide transparent access to user either at kernel level or in the form of an overloaded operator or networked object: e.g. A[103040304030] = 5.54;
- Useful in Tele-Immersion when visualizations get too big and it takes too long to load from disk- as in CIBRView case.

Collaborative Ubiquitous Computing Environments

Merging advanced computation, visualization and collaboration capabilities into Think Rooms

Maximally Co-Located Think Rooms, Project Rooms, War Rooms

- Olson & Olson – Michigan
- Study of 9 project rooms – software design, appliance design, organization design, sales response team
- Reported initial fear that co-located work may be distracting. Disadvantages of co-location was easily overcome by the benefits
- Comparison of war groups with norm showed performance of war group well above corporate average (doubling for software design case)
- Co-located and at a stage deemed appropriate for this kind of effort

Maximally Co-Located Think Rooms, Project Rooms, War Rooms

- Their time was not shared with other projects
- Fluidity of subgroup formation was rated very important
- Flip charts were persistent configurable records of their past meeting results. Charts could be moved anywhere
- Spatiality was important. Relationship of participants with artifacts in the space. Deictic references to them.
- Digital white boards saved time
- Small nearby cubicles to provide quiet but maintain accessibility to others

Lots of Related Work: http://www.research.microsoft.com/ierp/

- **Andersen Consulting CSTaR**: [ActiveEnvironments] [MusicFX] [HomeLab] [MagicWall]
- **AT&T Laboratories Cambridge**: [Overview] [ActiveBats]
- **ATR MIC:** [Creative Space]
- **U California, Berkeley:** [NotePals]
- **U California, San Diego:** [CVRR Lab]
- **Carnegie Mellon**: [Aura: Invisible Computing] [Spectrum Sharing]
- **Carnegie Mellon / Universität Karslruhe:** [Multimodal User Interfaces]
- **U Colorado, Boulder:** [Adaptive House]
- **FXPal:** [SmartSpaces]
- **Georgia Tech:** [Classroom 2000] [Future Computing Environments] [Domisilica]
- **GMD-IPSI:** [Ambiente]
- **IBM:** [Natural Interaction] [Pervasive Computing]
- **INRIA, Rhone-Alpes**: [PRIMA]
- **Universität Karslruhe :** [TecO: Ubicomp]
- **Universität Karslruhe / Carnegie Mellon :** [Multimodal User Interfaces]
- **Microsoft Research:** [EasyLiving] [User Modeling]
- **MIT:** [Oxygen] [House of the Future]
- **MIT AI Lab:** [Forest of Sensors] [Intelligent Room] [HAL]
- **MIT Media Lab:** [KidsRoom] [SmartRooms] [It/I] [VisMod] [TTT, Things That Think Consortium]
- **Nara Institute of Science & Technology:** [Knowledgeable Environment]
- **NIST**: [AirJava]
- **Rutgers, CAIP Center:** [Multimodal HCI] [Collaborative Information Processing] [Autodirective Microphone Arrays]
- **SRI**: [CHIC]
- **Stanford:** [Intelligent Workspaces] [Project Archimedes] [Responsive Workbench]
- **Trinity College Dublin:** [Intelligent Interfaces & Buildings]
- **UCLA**: [MSL - PTES]
- **USC**: [SCADDS]
- **VSAM (Visual Surveillance & Monitoring):** [Overview]
- **U Washington:** [Portolano]
- **U Wollongong:** [EACCS]
- **Xerox PARC:** [Digital Video Analysis] [PARCTAB] [Ubiquitous Computing] [ZombieBoard]

Distributed Coordinated Think Rooms: Questions & Issues

- What is appropriate Distributed Think Room work? Finding users.
- Electronically mediated collaboration is not good for tightly coupled work (Olson). [But global organizations seem to make it inevitable]
- How many people to a distributed Think Room? Classically 7 to a real room. So 7 to a site or 7 total?
- Patterns of work activity in distance Think Rooms – Sub-groups partitioned by locale. Concurrent loosely coupled work punctuated by tightly coupled synchronization.
- Identifying the most important aspects of co-location that need to be emulated in distributed Think Rooms for long term work.
- Merging Think Rooms with new computational, visualization, and networking technology. When you click on a terabyte document what does it launch?
- Timezone differences. Overcoming mis-matched circadian rhythms.
- Think Rooms are scare resources in a large organization. How can everyone get access?
- How do we measure the benefits?

Distributed Coordinated Think Rooms:
Thoughts on what needs to be supported based on our present understanding of collaborative work and non-distributed Think Rooms

- Persistence is necessary:
  - To overcome timezone differences
  - To overcome limited accessibility of Think Rooms
- Fluid sub-group formation.
  - Private asides between small sub group or individuals
- Spatial persistence of information artifacts.
- Private work spaces (cubicles) that also keep workers accessible.
- Comprehensive distribution, presentation and control of information artifacts via a whole range of input and output devices- e.g. data sets, visualizations, spreadsheets, documents, etc, viewable, controllable from PDAs, laptops, tiled displays, VR.

Distributed Coordinated Think Rooms:
Thoughts on technological configurations & requirements.
This is what we want to merge, somehow.


CAVE


AccessGrid

IDesk

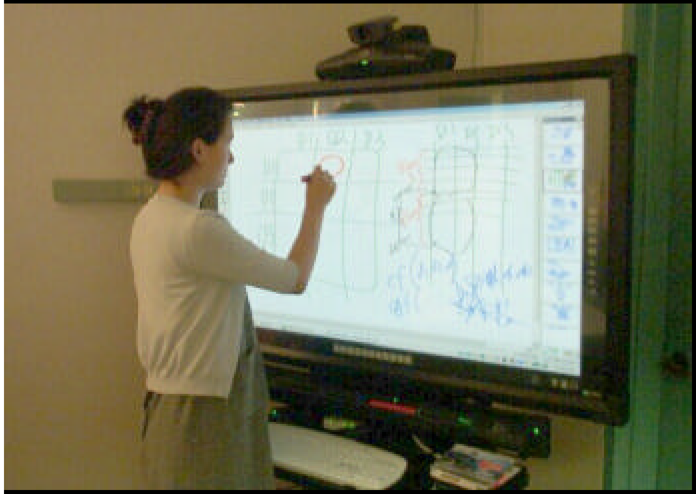
Plasma Touch Screen


AGAVE: Passive Stereo Wall

PDAs, Tablet PCs, Laptops

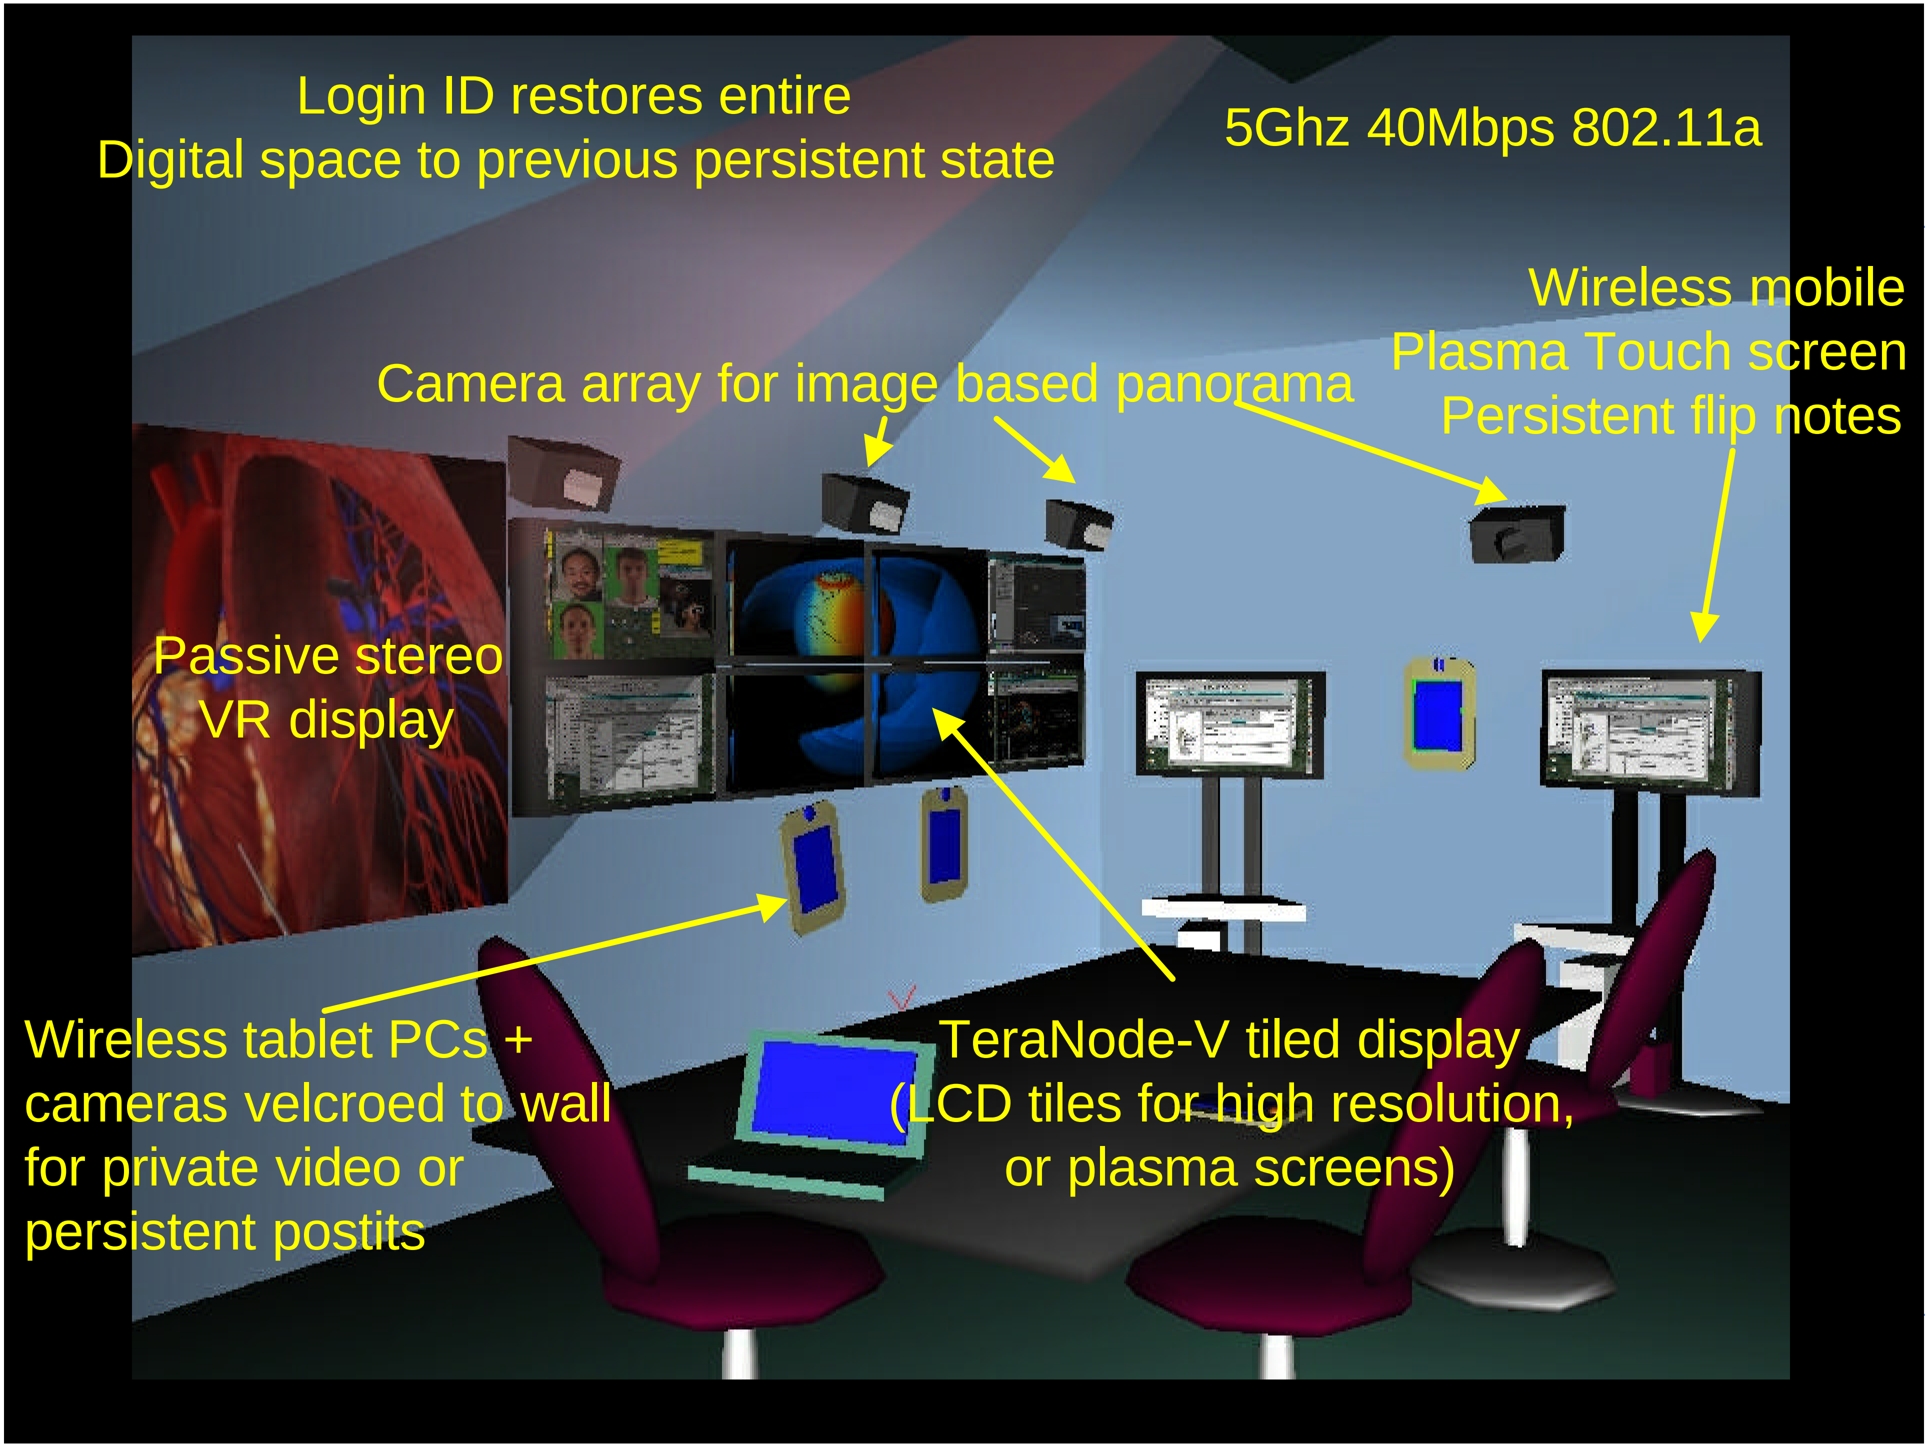

Login ID restores entire
Digital space to previous persistent state
5Ghz 40Mbps 802.11a
Wireless mobile
Plasma Touch screen
Persistent flip notes
Camera array for image based panorama
Passive stereo
VR display
Wireless tablet PCs +
cameras velcroed to wall
for private video or
persistent postits
TeraNode-V tiled display
(LCD tiles for high resolution,
or plasma screens)

Some Software Issues

- Interfaces for launching and controlling Supercomputing, AGAVE or Tiled Display applications from the Think Room
- Support opportunistic migration of data between media
  - E.g. I have jotted notes on my PDA, I want to drag and drop it to an AG screen for everyone to see.
- Develop smart filtering/translation systems for mediating collaborations between asymmetric display devices
  - E.g. for any given data set filters to limit the bandwidth to each device- based on constraints of the device, like- PDA, tablet computers, desktop computer, CAVE.
- Understanding strengths of each medium and building user interfaces and visualizations to take advantage of those strengths or adapt to their limitations
  - E.g. 3D collaborative interaction between a desktop computer using a mouse, and a CAVE using a 3D wand.
  - E.g. different types of visualizations for different devices.
- Long term collaboration – supporting persistence.
  - Need to expose the state of software so that it can be recorded and resumed. Easy for new tools. Harder for legacy tools.
- Need for security is obvious.

www.evl.uic.edu/cavern

spiff@evl.uic.edu